

## Research note 38

# Ethical AI Models of Code

Edward McDaid & Sarah McDaid

9 Apr 2025

**Many leading AI providers are embroiled in legal disputes around the use of other people's intellectual property in the creation of their models. At Zoea we completely avoid such problems in a simple but effective way.**

Every day seems to bring another news story about the owner of some large body of intellectual property suing one or more AI companies for (alleged) wholesale copyright infringement. These suits cover the complete spectrum of human creative output, from music and paintings to books and code.

What is common to all of these sorry cases is the modus operandi. AI models desperately need data and the more the better. In the race to dominate a growing and potentially huge AI industry, many of the players seem prepared to stretch their interpretation of the law to the limits of credulity.

Fair use is a common defence, although to be plausible this typically wouldn't involve using a particular work in its entirety - let alone many such works. It is also claimed that the resulting models do not contain verbatim copies of the source material. Given the nature of models as vast collections of numeric weights this is hard to prove either way. However, this can happen and the extent to which it occurs is currently disputed.

Some players even contest that any information which is accessible on the internet is fair game as training material for their models. Keen to cash in governments, regulators and professional bodies are happy to sit on their hands or even to actively facilitate questionable behaviour.

AI company leaders probably believe that they have no alternative. They see the risk of failing to carve a sizeable niche in the current AI bubble as an existential threat - so they basically stand to lose everything. In the interim, AI companies have deep pockets when it comes to lawyers and litigation of this sort often takes years. Some of the plaintiffs will no doubt cut deals, simply to avoid being bankrupted by the process. In the end, so their thinking goes, the AI industry will simply become too big to be held to account. As with data protection, any resulting sanctions will just be a relatively trivial part of the cost of doing business.

At the same time, there are also actors in other countries where attitudes to intellectual property are more relaxed. If ethical considerations are allowed to hinder technology companies in the developed world then they will quickly fall behind competitors who aren't so encumbered. Aside from threatening the economic stability of the Free World, all manner of bad outcomes are sure to ensue.

Fortunately, it doesn't have to be this way. As it happens, it is entirely possible to build AI systems while still respecting the rights of intellectual property owners. Indeed, this is what we do at Zoea. To understand how we manage this it is first useful to examine why the problem exists in the first place.

The root cause of all these problems may seem somewhat surprising. Fundamentally, it is down to the lack of any interest - on the part of most of the AI community - in the knowledge that is encoded into the systems that they build. This kind of 'shut up and calculate' attitude leads to things like the Copenhagen interpretation of quantum mechanics. That is, being satisfied with building a model which is useful, without the slightest interest in how or why it works.

The widespread adoption of neural networks over the last couple of decades has had the same sort of chilling impact on AI. Not only did it effectively kill off any curiosity about how AI systems work but it also made it more or less impossible to find out. This kind of mindset is directly responsible for many of the problems facing AI. These include the inability to create transparent or explainable AI, issues of bias and a total lack of security. It also makes it impossible to foresee, diagnose and fix unintended consequences like hallucinations.

With mainstream AI, huge amounts of data is fed into the system at training time, like a giant sausage machine. Nobody understands what the resulting model contains or how it is used to produce solutions. Indeed, nobody really cares - so long as it works well enough, most of the time. The share price is the key driver for innovation.

At Zoea we also build models but the key difference is that we know exactly what knowledge we use, how it is represented and also how it is processed. This is because knowledge in Zoea is represented explicitly. The models that we build currently relate to code but there is no reason why our approaches couldn't be applied to other sorts of knowledge as well.

The explicit representation of knowledge has a long history in AI - although it has largely fallen out of favour since the 1990s. This is because it is a common fallacy to conflate explicit knowledge representation with the manual elicitation and encoding of knowledge. However, these two approaches are not the same thing, nor even mutually inter-dependent. For example, it is entirely possible to automatically extract knowledge which is also represented explicitly. This is what we do at Zoea and our approach addresses all of the alleged problems associated with explicit knowledge including brittleness, bias and overfitting - while retaining all of the benefits of transparency and explainability.

Unlike mainstream AI such as LLMs, the models that we create at Zoea are also unquestionably ethical. This is because we extract and store no intellectual property when we analyse an example source code program. Instead, all of the information we are interested in is generic and could apply to millions of other different pieces of code. For example, in order to produce instruction subsets we simply list the names of the different programming language instructions that are used in each program. This information cannot belong to anyone as exactly the same list of instructions is used in very many other programs, written by different people. Were this not the case then virtually all code would already be breaking someone else's copyright. The same argument also holds for information describing which instructions invoke other instructions. This is the Zoea instruction digram model.

The models that Zoea creates do not contain a single line of anyone else's code. Yet we are able to describe all of the explicit and much of the tacit knowledge that people use when writing software. The only additional overhead required is a little more effort to

understand and identify the information that we want to extract. All of this information corresponds to a tiny fraction of the size of the training set and is effectively limited to summary statistics. By anyone's definition this is a real example of fair use.

Learn more at [\*\*zoea.co.uk\*\*](https://zoea.co.uk)